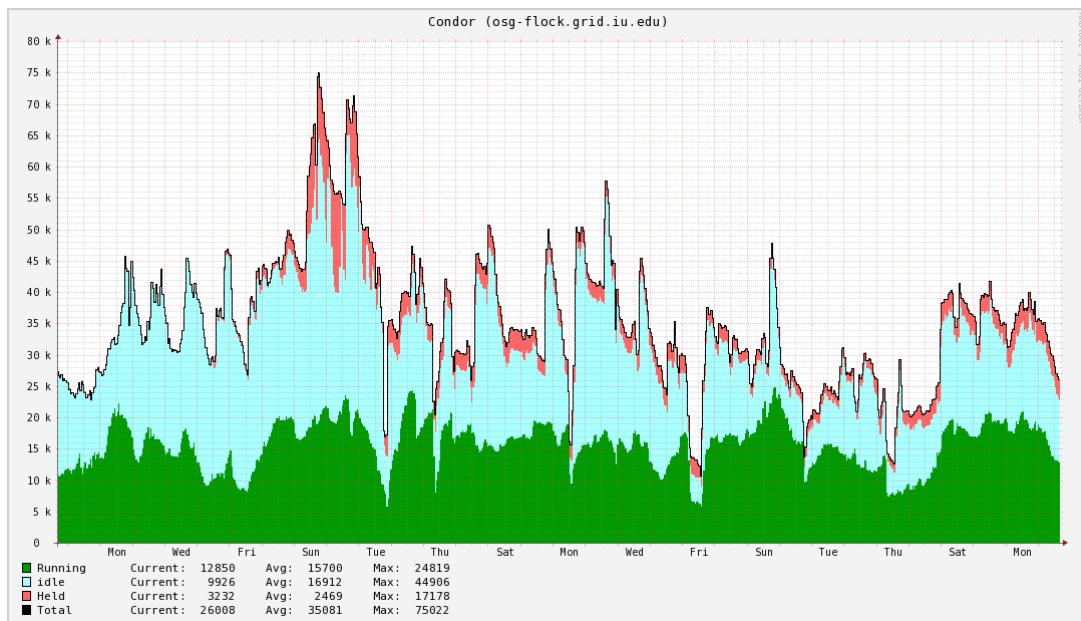


Less Demand?



User	AcctGroup	R	I	H
<hr/>				
Q1@xd-login.opensciencegrid.org:				
aantipov	xsedehigh	9	0	0
dsschult	xsedeelow	14	0	0
snowball	opportunistic	202	1	0
strolog	opportunistic	1096	175	1788
yx59	opportunistic	6	0	6
<hr/>				
Q3@xd-login.opensciencegrid.org:				
jauvinen	opportunistic	11	0	0
weiyao-ke	opportunistic	0	0	1
<hr/>				
Q4@xd-login.opensciencegrid.org:				
yx5	opportunistic	8304	10135	0
<hr/>				
duke-login.osgconnect.net:				
lincolnb		0	1	0
<hr/>				
glide.bakerlab.org:				
possu		0	0	3
<hr/>				
login01.osgconnect.net:				
vo.uc3.bburnett	opport	0	1	0
vo.uc3.dbala	opport	0	2	0
vo.uc3.dschatz	opport	1	0	0
vo.uc3.fluco	opport	0	3	0
vo.uc3.hbhat	opport	20	3	0
vo.uc3.jenkins	opport	0	1	63
vo.uc3.jtqv84	opport	0	257	0
vo.uc3.leeping	opport	1769	6	1109
vo.uc3.nathankaib	opport	115	2	2
vo.uc3.oanacarja	opport	26	1	0
vo.uc3.seanbc	opport	0	1	0
vo.uc3.tzhao	opport	0	2	0
vo.uc3.vjneal	opport	0	1	0
vo.uc3.yadunand	opport	0	0	1
<hr/>				
login02.osgconnect.net:				
vo.uc3.jbvono	opport	0	0	1
<hr/>				
xd-login.opensciencegrid.org:				
psiders	xsedehigh	53	0	0

XDMoD update

- Background: XDMoD did not have good OSG data due to us sending “summary” jobs to the XDCDB
- We are now sending raw job records directly to XDMoD
 - Only records from `xd-login.opensciencegrid.org`, with project names starting with TG-
- The data will eventually be merged into the main XDMoD database

GPUs and large memory jobs

Large memory works

GPUs works, but we still have to advertise GPU model and attributes

The screenshot shows a web browser window for the 'Support : OSG Helpdesk - Google Chrome'. The address bar shows the URL: support.opensciencegrid.org/support/home. The page has a header with the 'opensciencegrid help desk' logo and navigation links for Home, Solutions, and Forums. On the left, there's a search bar with placeholder text 'Enter your search term here...' and a 'SEARCH' button. Below it is a 'How can we help you today?' section with a search bar. The main content area is divided into two main sections: 'Knowledge base' and 'Community forums'.

Knowledge base:

- Overview**
 - Getting Started (4)
 - A welcome from the User Support Team
 - Communicate with us via Twitter
 - Is high throughput computing for you?
 - Acknowledging the Open Science Grid
 - OSG XSEDE Users (1)
 - OSG XSEDE User Guide
- OSG Connect User Guide**
 - Getting Started with OSG Connect (5)
 - Registration and Login for OSG Connect
 - OSG Connect Quickstart
 - Start a Project with OSG Connect
 - The 'tutorial' command
 - Compiling Applications for OSG Connect
 - Submitting Jobs (4)
 - Steer your jobs with HTCondor job requirements
 - Large Memory Jobs
 - GPU Jobs
 - Multicore Jobs
 - Software access (4)
 - Accessing Software using Distributed Environment Modules
 - Software transfer via HTCondor or HTTP
 - Software modules catalog
 - Requesting a software installation
 - Transferring Data to OSG Connect (2)
 - Using scp to transfer files
 - Data transfer with Globus

50k running jobs scaling test

ET has asked us to show that the infrastructure can handle 50k running jobs

Edgar/Bo getting sleeper pools lined up, which we will add on top of production workload.